ChatGPT

Abstract

This paper examines a recent incident in which a digital research repository, Figshare, removed an entire portfolio of thirty research papers after a single new submission – "Toward a World Without War: OntoMotoOS and Philosophical Reinsurance" – was flagged for containing a ChatGPT logo. The incident is analyzed as a case study at the intersection of institutional policy enforcement and scholarly norms of transparency, especially in the context of human–AI co-created research. We begin by introducing **OntoMotoOS**, a co-creative philosophical framework that integrates human and AI collaboration in knowledge-making. We then present the facts of the Figshare takedown with balanced detail: a **platform policy** against AI-generated content led to the simultaneous removal of 29 previously accepted papers once the 30th paper triggered scrutiny due to its AI disclosure logo [1]. Through a theoretical lens, we explore the **philosophical tension between institutional consistency and scholarly transparency**. Institutions like publishing platforms prioritize consistent rule application and content integrity, whereas scholars prioritize truthfulness about methods and tools [2]. We discuss how this clash highlights deeper questions of policy scope, platform design, and academic values in the age of AI-assisted research. While critically examining the scenario, we maintain a respectful tone, acknowledging the rationale on both sides. In conclusion, we reflect on how the episode, despite its discouraging outcome, contributed to the author's ongoing philosophical inquiry. In a gesture of constructive dialogue, the author expresses appreciation toward Figshare for prompting these reflections – a reminder that intellectual progress can arise even from conflict.

Introduction

Advances in generative artificial intelligence are transforming scholarly practice, enabling new forms of human–AI co-authorship and **co-creative research**. This development has prompted active debate on how academic norms should evolve – particularly concerning transparency about AI assistance and the enforcement of integrity policies [3] [4]. As researchers experiment with large language models (LLMs) in writing and analysis, academic institutions and platforms face the challenge of updating guidelines to ensure responsible use of AI tools [5] [6]. A growing consensus among publishers is that **AI tools may be used but must be disclosed**, and that human authorship and accountability remain paramount [5] [7]. However, this push for transparency can sometimes collide with **institutional policies** designed to guard against unethical practices, such as wholly machine-generated papers. The case under study exemplifies this collision and its broader implications for scholarly communication in the digital age.

In mid-2025, Figshare – a popular open repository for research outputs – **removed 30 of an author's papers** after flagging one new submission for a policy violation related to AI-generated content [1]. The flagged paper, titled "Toward a World Without War: OntoMotoOS and Philosophical Reinsurance," was an academic exploration of global peacebuilding, notable for explicitly involving an AI tool (OpenAI's ChatGPT) as part of its methodology. The author had placed a **ChatGPT logo** on the paper's title page as a mark of **scholarly transparency** – a disclosure of the AI's contributory role. From the author's perspective, this logo served to openly acknowledge human–AI collaboration in the research, aligning with emerging ethical norms that call for researchers to **disclose the use of AI tools in their work** [8] [6]. However, to Figshare's automated enforcement system (and content moderators), the presence of the AI logo was interpreted as a sign that the paper "consists primarily of content generated by artificial intelligence," a category of content disallowed by Figshare's terms of service [1] [9]. The immediate consequence was the **disabling of all 30 DOIs (Digital Object Identifiers)** associated with the author's

Figshare account, effectively removing access to 29 previously published items alongside the new paper [10] [11].

This incident raises critical questions at the heart of **digital scholarship and platform governance**. How should platforms balance a need for **consistent policy enforcement** with the academic imperative for **transparency and openness** about research processes? What does this conflict reveal about the design of scholarly repositories in an era when AI is increasingly intertwined with human creativity? And how might conceptual frameworks for human–AI collaboration – such as the author's own **OntoMotoOS** philosophy – inform a more nuanced resolution to such conflicts? In what follows, we provide a theoretical background on the OntoMotoOS framework and its ethos of co-creation, describe the Figshare case in detail, and then delve into a discussion of the philosophical and institutional implications. The analysis will highlight the tension between **institutional consistency** (strict adherence to platform policies for the sake of integrity and order) and **scholarly transparency** (the researcher's duty to truthfully represent the role of AI in knowledge production). We aim to examine both sides of this tension respectfully, recognizing the legitimate principles involved on each side. Finally, we conclude with broader reflections and a note of appreciation, suggesting that this episode, however unfortunate for the author, has yielded valuable insights for the academic community. The incident serves as a timely microcosm of the challenges and opportunities in aligning **academic values** with **platform design** and policy in the age of AI-enhanced scholarship.

Theoretical Background: OntoMotoOS and Co-Creative Scholarship

OntoMotoOS (Ontology–Motor Operating System) is a conceptual framework developed as a "philosophical operating system" for human–AI collaboration [12]. In essence, OntoMotoOS proposes a **co-creative approach** in which human reasoning and artificial intelligence are integrated into a single iterative process of knowledge generation, decision-making, and ethical deliberation. It is described by its creator as a "mesh-based, recursive meta-OS for unifying AI, ethics, [and] ontology", emphasizing that the system operates across different layers of understanding to merge abstract values with concrete computational operations [13] [14]. Rather than treating AI as a mere tool or an autonomous agent, OntoMotoOS envisions AI systems as **"collective co-thinkers"** working alongside humans in a transparent, interactive network of problem-solving. This framework draws from multiple domains – philosophy, computer science, systems theory – to outline how recursive alignment between human values and AI behavior might be achieved through continuous feedback loops and shared representations of knowledge.

One motivating application of OntoMotoOS, as explored in the removed Figshare paper, is the prevention of war and violent conflict [15] [16]. The concept of **"philosophical reinsurance"** introduced in that work refers to an unconventional safeguard at the level of fundamental values and norms [17] [18]. By analogy to financial reinsurance (insurance for insurers), philosophical reinsurance is a higher-order safety net for civilization – a way to insure humanity's future against catastrophic moral failures (such as descent into war or tyranny) by embedding our deepest ethical principles into the architecture of our decision-making systems [19] [20]. OntoMotoOS is proposed as the operational embodiment of this idea: it seeks to instantiate core humane values (like justice, empathy, and rational conflict resolution) into a working, evolvable **meta-system** that can simulate outcomes, detect ethical risks, and recommend interventions before crises escalate [21] [22]. In simpler terms, OntoMotoOS provides a "sandbox" where ideals (e.g. a world without war) can be rigorously tested against real-world complexities via AI-driven simulation and human oversight [23] [16]. The framework's **co-creative nature** is crucial: humans input goals, values, and contextual knowledge, while AI components contribute by processing vast data, running scenario

analyses, and even offering creative suggestions – all within guardrails defined by philosophical and ethical constraints.

Underpinning the OntoMotoOS philosophy is a commitment to **transparency and reflexivity** in human–AI interactions. The system is recursive, meaning outcomes loop back into re-examining the initial values and assumptions, fostering an ongoing dialogue between abstract principles and concrete results [24] [25]. This stands in contrast to treating AI as an opaque oracle or a one-off tool; instead, AI's role is explicit and open to scrutiny at each stage of reasoning. It is within this intellectual context that the author's inclusion of a **ChatGPT logo** in a research manuscript can be understood: it is an **artifact of the OntoMotoOS ethos**, symbolizing that AI (in this case, ChatGPT) was a visible participant in the thought process. The presence of the logo signaled an adherence to the emerging academic norm that any use of AI in producing results or text should be **fully disclosed to readers** [26] [6]. This norm is increasingly reflected in publisher policies and guidelines, which insist on disclosure of AI assistance in methods sections or acknowledgments in order to maintain trust in the scholarly record [27] [7]. From a theoretical standpoint, such transparency is not only a matter of honesty but also aligns with the co-creative ideal: if human and AI are co-authors of an idea, then both presences should be acknowledged in the intellectual lineage of that idea.

However, the OntoMotoOS vision of seamless human–AI cooperation also brings to light a potential friction with existing **institutional structures**. Traditional academic and publishing institutions operate on principles of **consistency, quality control, and clear attribution of intellectual responsibility**. The introduction of AI as a quasi collaborator blurs long-standing categories – for instance, who or what qualifies as an author, and how do we evaluate work that is partially machine-generated? Institutional policies have begun to take shape in response to these questions: for example, many journals and conferences now explicitly forbid listing AI systems as authors, and require authors to attest that they, not an AI, take responsibility for the content [8] [6]. Some platforms outright ban submissions that are "primarily AI-generated" on the grounds that such content might be nonsensical, plagiarized, or lacking the accountability that human scholarship demands [28]. Figshare's policies, in particular, include removal rights for any content that "consists primarily of content generated by artificial intelligence or other automated systems, rather than by the user." This creates an inherent tension when an author intentionally marks their work with evidence of AI involvement: the author is signaling intellectual honesty and methodological clarity, yet the platform's safeguards may interpret that very signal as a red flag for policy violation. The **OntoMotoOS framework**, advocating co-creation "without erasure" of either human or AI contributions [29], thus runs up against a legacy model of academic gatekeeping which has not fully adapted to nuanced human–AI partnerships. This theoretical backdrop sets the stage for understanding the specifics of the Figshare case and its broader meaning.

Case Description: The Figshare Takedown of AI-Tagged Research

Background of the Incident. In September 2025, the author of this paper (an independent researcher and creator of the OntoMotoOS framework) uploaded a new working paper to Figshare, an open-access repository that issues DOIs for shared research outputs. This paper – "Toward a World Without War: OntoMotoOS and Philosophical Reinsurance" – was the latest in a series of publications developing the OntoMotoOS concept and examining its applications to global governance and peace. The manuscript was notable for its forthright disclosure of AI assistance: the title page or appendix included the logo of **ChatGPT**, indicating that the large language model had been used as a collaborative tool in the writing or research process. According to the author, the inclusion of the ChatGPT emblem was a deliberate act of **academic transparency**, meant to inform readers that AI had a role in generating ideas or text – a practice aligned with the broader ethical call for openness about AI use in scholarship [30] [8]. The author had already disseminated this work through other channels (such as SSRN and PhilArchive)

without incident, suggesting that the content itself had academic merit and was not simply AI gibberish [31].

Platform Policy and Automated Flagging. Shortly after uploading, the paper was **disabled by Figshare staff or algorithms**, citing a "violation of policy against AI-generated content." Figshare's content guidelines at the time explicitly prohibited works that appeared to be predominantly machine-generated, reflecting a precautionary stance against the influx of low-quality or uncredited AI outputs on scholarly platforms [1]. The presence of the ChatGPT logo seemingly triggered this policy: in the eyes of the repository's enforcement mechanism, the logo was interpreted as an admission that the content might not be primarily the user's original work. In other words, a badge meant to denote **AI usage transparency** was read as evidence of **AI-generated content** subject to removal. Figshare's response was sweeping. Instead of removing only the single offending item, **all items associated with the author's account were taken down simultaneously** [10]. In total, 30 DOIs (each corresponding to a separate paper or dataset) were **deactivated**, effectively erasing the author's entire body of work from Figshare's public interface. Even the author's profile page was removed, leaving no trace on Figshare's site of the contributions that had previously been available [11].

This blanket action was ostensibly taken in the name of **consistent enforcement**. From Figshare's perspective, allowing one flagged item to remain or only a partial rollback could set a precedent undermining their rules [32]. The rationale might be summarized as follows: if an exception were granted for a paper visibly marked as AI-assisted (even if the paper's content was scholarly), it could encourage others to upload fully AI-generated works under the guise of "AI-assisted," thereby flooding the platform with potentially unreliable content. Figshare's policy apparatus did not (at that time) have a nuanced way to distinguish a legitimate research paper that discusses or uses AI from a paper that is actually written by AI with minimal human input [31]. Thus, to **avoid any appearance of inconsistency or favoritism**, the platform opted to remove all of the author's content, treating the entire portfolio as suspect once the violation threshold was crossed. In effect, the author was deplatformed from Figshare overnight, losing the repository links to 29 other works that had not been individually challenged on any grounds of validity or originality. Those works included prior installments of the OntoMotoOS series and related essays on AI ethics and governance, which had been available and even cited in the interim. The sudden disappearance of these resources from Figshare raised alarm not only for the author but for any scholars who might have been accessing or referencing those DOIs.

Author's Response and Perspective. The author faced a stark choice in response to Figshare's takedown: either **remove the ChatGPT logo and attempt to reinstate the work**, or uphold the principle of transparency and accept the loss of the Figshare entries. The author chose the latter, doubling down on the conviction that the **logo "must stay" as a matter of scholarly principle** [33]. In a contemporaneous reflection, the author wrote: "Removing the logo is non-negotiable in my practice of scholarly transparency", emphasizing that future academia ought to acknowledge AI contributions rather than conceal them [34]. In the author's view, yielding to the platform's demand (implied by the policy) to eliminate visible signs of AI assistance would have meant participating in a **ritual of hiding the truth** about how the research was conducted. It would also, arguably, send the wrong signal – that using AI is something shameful or to be kept in the shadows, which contradicts the author's co-creative philosophy. The author instead sought a kind of **principled resolution**: Figshare would maintain its stance, and the author would maintain theirs, resulting in a parting of ways. Specifically, the author framed it as a "win-win resolution" in which each side preserves its integrity: **Figshare** demonstrates its strict policy compliance by applying the rules uniformly (removing all 30 DOIs, no exceptions) and **the author** demonstrates fidelity to academic transparency by retaining the AI attribution on the work and continuing to share it through alternative platforms [35]. Indeed, the paper and the broader set of works remain accessible via other channels (SSRN eJournal series, PhilPeople profile, etc.), and thus the research community can still find them – just not on Figshare [36] [37].

Outcome and Immediate Implications. The immediate outcome is that, as of this writing, the author's contributions have vanished from Figshare's public repository. Intriguingly, however, a kind of **scholarly trace** persists: the DOIs that were registered for those items are still indexed in systems like ORCID and Google Scholar, albeit pointing to inactive Figshare pages [38] [39]. This means that the metadata (titles, authors, etc.) of the works remain part of the permanent scholarly record, even though Figshare no longer hosts their content. As the author noted, "Figshare may disable a link, but it cannot delete the fact of the work" [40]. This dichotomy between the persistence of scholarly metadata and the ephemerality of platform access highlights a unique aspect of the case: it underscores how **digital scholarship is simultaneously robust and fragile**. On one hand, DOI infrastructure and cross-platform dissemination ensured that the work was not completely lost; on the other hand, the centrality of platforms like Figshare for visibility means that removal still inflicts substantial damage on the work's accessibility and the author's academic footprint. The author expressed **regret and disappointment** at the outcome, calling it "profoundly discouraging" to see an entire body of work effectively erased after one perceived misstep [41]. The countless hours of research appeared wasted in that venue, illustrating a personal cost to rigid enforcement. Nonetheless, the author also derived a philosophical lesson from the event – which leads into our discussion: perhaps such conflicts are inevitable in this transitional era of AI in academia, and they force us to confront what principles we truly value and how to build systems that honor those principles without collateral damage.

Discussion: Institutional Consistency vs. Scholarly Transparency

The Figshare–OntoMotoOS incident serves as a vivid microcosm of a broader **philosophical tension** in the digital age: the conflict between rule-based institutional consistency and context-sensitive scholarly transparency. At its core, this tension is about **trust** – but different institutions cultivate trust in different ways. **Platforms** and repositories build trust by ensuring content quality, reliability, and adherence to norms, often through clear rules and uniform enforcement. **Researchers** build trust by demonstrating honesty, rigor, and openness in their methods, which includes disclosing the use of new tools like AI. When these modes of trust-building align, the system functions smoothly. In our case, however, they came into conflict, prompting us to examine the values and assumptions on each side.

The Institutional Perspective: Integrity Through Consistent Policy

From the institutional standpoint (here embodied by Figshare), the paramount value was **integrity of the repository** – ensuring that Figshare does not become a dumping ground for low-quality or misleading content. The policy against "AI-generated content" can be understood as a protective measure against a foreseeable problem: large language models make it trivially easy to generate voluminous text that mimics academic writing but lacks genuine scholarship. If left unchecked, such content could **undermine the credibility** of open research repositories. By adopting a zero-tolerance rule (no primarily AI-written uploads) and enforcing it **consistently**, Figshare signals to all users that it takes quality control seriously [42] [32]. In ethical terms, this approach resonates with **rule-utilitarianism** or **deontological ethics** – the idea that following a general rule (e.g., "do not host AI-generated papers") produces the best overall outcome or upholds a duty, even if in individual cases the rule might seem too strict. The logic is that if one clear violation is overlooked, the credibility of the entire platform's review process could be called into question. Consistency is equated with fairness (no special treatment) and with predictability (users know the rules and consequences).

In the Figshare case, the decision to remove **all 30 items** rather than just the single flagged paper exemplifies the lengths to which the platform went to maintain **policy consistency**. One might argue that this response was disproportionately harsh; yet from Figshare's perspective, it may have seemed necessary to avoid any ambiguity. A partial measure (e.g., removing or editing the logo in one PDF) might

implicitly condone similar content in other uploads. Moreover, it might have been technically simpler to deactivate the whole account pending investigation rather than surgically remove one item – pointing to a **platform design** issue where moderation tools operate at the account level. This reveals a critical facet of platform governance: blunt instruments are often used when nuanced ones are unavailable. The case highlights how **platform design choices** (such as the lack of a feature to label content as "AI-assisted but not AI-generated") can force binary outcomes – keep or delete, with little room in between. Figshare had "no mechanism to differentiate between research about AI and AI-written content" [31], so its enforcement could not help but be over-inclusive.

We can also view Figshare's actions through the lens of **institutional risk management**. In a rapidly evolving area like AI-generated text, erring on the side of caution is common. Making an example by strict enforcement can deter would-be violators. It sends a message aligning with what Figshare likely perceives as its responsibility to the academic community: to maintain a repository where entries can be presumed genuine and authored by accountable individuals. The downside, of course, is that such an approach **sacrifices nuance for certainty**. In this incident, the "collateral damage" was the removal of legitimate human-authored research (indeed, research that was peer-recognized on other platforms) because it carried an AI marker. The institutional perspective would justify this by claiming that rules, to be effective, must sometimes be inflexible. There is an echo here of the classic philosophical debate between **justice as uniform application of law** versus **equity as case-by-case fairness** – Figshare opted firmly for the former.

The Scholarly Perspective: Integrity Through Transparency and Truthfulness

From the scholar's standpoint (the author and many in the academic community), **integrity in research** is bound up with **transparency** and truth-telling about one's methods and sources. The inclusion of the ChatGPT logo can be seen as an embodiment of the scientific norm of **methodological transparency**: just as one would report using a particular laboratory instrument or dataset, the author reported using an AI assistant for idea generation or writing. Far from undermining trust, this disclosure enhances trust by allowing peers to understand the context in which the work was produced [6] [43]. Opaque or undisclosed use of AI, by contrast, could be viewed as a form of potential plagiarism or misrepresentation (if text is copied from a model without credit) [44] [6]. Thus, the author's refusal to remove the logo was a stance grounded in **academic ethics**: it was about preserving the truthful provenance of the work and affirming a commitment to the value of openness. In a time when many institutions are grappling with an "AI trust deficit," individual scholars taking initiative to be transparent might help set positive norms [6].

Philosophically, the scholar's stance aligns with a view of **ethics of care or virtue ethics** in academia – emphasizing honesty, courage (to admit using a non-traditional tool), and a form of humility about the research process (acknowledging that one had help from an AI). It also reflects the principle that **ends do not justify means** if the means involve suppressing truth. The goal of having one's work widely accessible on Figshare, while desirable, was deemed not worth achieving through the "dishonest" act (from the author's perspective) of hiding the AI involvement. In other words, the **scholarly transparency principle** was ranked higher than the practical benefit of compliance with a rule that conflicted with that principle. This is reminiscent of Kantian ethics: one must not lie or mislead, even if honesty has inconvenient consequences. Interestingly, it is also reminiscent of the core ethos of science as articulated by Robert Merton's norms (Merton, 1942) – particularly communism (open sharing of knowledge) and organized skepticism. Here, "organized skepticism" could extend to being skeptical of one's own processes enough to disclose them for scrutiny.

Another facet of the scholarly perspective is the **value of context**. Scholars argue that not all AI-involved works are equal: there is a meaningful distinction between plagiarized or purely machine-generated text and legitimate research that uses AI as a tool. The author clearly situates their work in the latter category

– indeed the content of the paper was about AI's role in society, and the AI logo was contextually appropriate. A human researcher guided and curated the output, an important detail for assessing the work's credibility. By treating all AI-marked content as suspect, the platform's policy failed to consider context and proportionality. This is akin to a **philosophical critique of algorithmic governance**: rigid rules can't capture context, whereas human judgment, which scholars rely on in peer review, typically does. The scholar might ask: shouldn't there be room for academic judgment in deciding what constitutes acceptable use of AI, rather than a blanket ban? The absence of such discretion in Figshare's approach points to a gap between **academic values** and platform moderation practices.

The Collision and Its Broader Implications

When we analyze the incident, neither side's perspective is without merit – indeed, that is why this is a hard case. The clash illuminates several broader implications for the future of academic work:

1. Redefining Policy Scope for AI in Research: The case suggests that institutions may need to refine the scope of their AI-related policies. A binary distinction (AI-generated vs. not) is too crude. As AI becomes a common assistant in writing and data analysis, policies must distinguish between responsible use with disclosure and irresponsible use or wholesale generation. Some publishers have begun to articulate such nuanced policies – for example, by requiring authors to detail how AI was used, to ensure humans remain accountable for the content [45] [43]. Repositories like Figshare might consider adopting a similar approach: instead of banning "AI-generated content" outright, they could require a **declaration of AI assistance** and perhaps an additional layer of verification for such works. This could involve an **AI usage statement** as part of metadata, which curators or readers can review. By widening the policy scope to acknowledge degrees of AI involvement, platforms could enforce quality standards without automatically penalizing transparency. In short, a shift from prohibition to monitoring and documenting AI use could be more future-proof and fair.

2. Platform Design and the Need for Human-in-the-Loop: The incident highlights limitations in platform design regarding content moderation. Automated flagging based on an image (the AI logo) can catch a broad class of cases but will inevitably generate false positives (e.g., an article about AI that legitimately shows an AI logo). This calls for a more **human-in-the-loop approach** in moderation when AI is involved. Instead of immediate mass removal, a platform could implement an alert and review system: notify the author and/or a human moderator when AI content is detected, then allow a period for explanation or adjustment. Perhaps the author could have been given a chance to explain that the content was original and the logo was for disclosure purposes. The absence of an appeals process or dialogue is notable in this case – a design choice that prioritizes efficiency over deliberation. Philosophically, this ties into debates on **procedural justice**: stakeholders should have voice and explanation in decisions that affect them. Designing platforms that incorporate a dialogic process (especially for complex cases like AI-assisted work) could lead to outcomes that uphold rules and context-sensitive fairness. This also relates to **academic freedom** in the digital context – scholars might be more willing to embrace open platforms if they trust that those platforms won't abruptly censor or remove content without due consideration.

3. Academic Values and the Culture of Openness: The scenario opens a discussion on how strongly the academic community values openness and transparency, and whether those values are adequately protected. There is an argument to be made that **platforms serving scholars should align with scholarly norms**, not just impose their own. If transparency about AI usage is increasingly seen as an ethical requirement in research [6] [7], then platforms should arguably reward rather than punish it. This might mean developing better screening to differentiate fraudulent AI-generated papers (which likely come with other telltale signs such as lack of coherent references or plagiarism) from genuine research that simply uses AI. It could also mean engaging with researchers to update community

guidelines collaboratively. The academic community might consider establishing **shared standards for AI disclosures** such that platforms can easily verify and accept them. For instance, a standardized statement like: "This paper was drafted with the assistance of GPT-4. The authors have reviewed and edited all AI-generated text and take full responsibility for the content." If Figshare saw such a statement in the metadata or text, perhaps it could treat the item differently than one where an AI logo appears with no context. Thus, a conversation between scholars, platform providers, and organizations like COPE (Committee on Publication Ethics) could help harmonize the goals of integrity and transparency.

4. Preservation vs. Ephemerality in Scholarly Communication: An underlying institutional question highlighted by the takedown is: What are the obligations of a platform that mints DOIs and hosts content, in terms of preservation and continuity? The fact that DOIs were issued and then "orphaned" by removal might be viewed as troubling from a research integrity standpoint. DOIs are meant to provide a persistent link to information. If platforms can void them unilaterally, this might undermine confidence in the scholarly record. Some commentators might argue that once a DOI-bearing item is published (especially if it's not illegal or patently fraudulent), there should be strong hesitations against removal – perhaps preferring measures like "flagging" with warnings, or updating with a notice, rather than complete erasure. The case at hand indicates that current infrastructure did not prevent rapid removal, but at least the metadata survived externally. This points to the importance of **redundancy in knowledge dissemination**: relying on a single platform can be risky for scholars. Open science proponents might take this as a lesson to always deposit work in multiple repositories or use decentralized preservation (like LOCKSS – Lots of Copies Keeps Stuff Safe). Institutionally, there may be a need for clearer guidelines on what justifies retraction or removal in non-peer-reviewed repositories, akin to retraction policies in journals.

5. Ethical AI Collaboration as a New Frontier: Finally, the incident underscores a more philosophical implication: as we integrate AI into our cognitive processes, we will encounter new forms of **principle collisions**. The author pointed out that ironically, the conflict became a "real-world case study" of the very ideas in OntoMotoOS, which seeks "co-creative resolution without erasure." [29] In a sense, the platform "erased" content in order to enforce rules, rather than finding a way to resolve the underlying issue (the presence of AI in authorship) collaboratively. The challenge moving forward is whether our institutions can evolve co-creative solutions where human norms and AI's contributions find a productive equilibrium. For example, could AI itself assist in moderation by assessing whether a text is mostly human-written vs AI-written? Could there be AI auditors that help platforms enforce policies more intelligently, reducing false positives? And at an even higher level, is it possible to design systems (mirroring OntoMotoOS principles) where conflicts between individual scholarly practice and institutional rules trigger a constructive dialogue to improve both the platform and the practice? This calls for **philosophical reinsurance** of a sort within academia itself: backup measures and meta-level frameworks to ensure that as new technologies disrupt norms, we don't default to damaging reactions (like mass deletions), but rather engage in reflective equilibrium to adjust our values and systems.

In sum, the discussion reveals no villains, only trade-offs. **Institutional consistency** is crucial for maintaining a trusted archive of knowledge – scholars and readers rely on repositories to uphold certain standards. Yet **scholarly transparency** is equally crucial for the credibility of research – science advances by exposing, not obscuring, the conditions of discovery. The current case shows how easily one can undercut the other if implemented without communication. The broader implication is a call for **rethinking the design and governance of scholarly platforms**. The goal should be to create an environment where an honest researcher is never placed in the position of having to choose between telling the truth about their methods and having their work disseminated. Achieving this will likely require updates to policy (more granularity and flexibility), enhancements in platform technology (better AI detection tools and moderation interfaces), and ongoing dialogue about academic norms as they pertain to AI.

Conclusion

The Figshare incident involving the OntoMotoOS research papers is more than an isolated disagreement over a logo – it is a **symbolic episode** highlighting the growing pains of academia in the age of AI. We have analyzed how the removal of 30 papers following an AI disclosure emblem brought into sharp relief the **competing principles** of a platform's rule enforcement and a scholar's commitment to transparency. This analysis has situated the case within a larger context of evolving publisher guidelines, ethical frameworks for human–AI collaboration, and the fundamental values that underlie academic work. On one side, we see the importance of **institutional consistency**, ensuring that emerging threats like unchecked AI-generated content do not erode the integrity of scholarly communication. On the other side, we see the imperative of **scholarly transparency**, upholding honesty and openness as non-negotiable pillars of research integrity [6]. The clash between these sides in this instance was stark, but it also offers a learning opportunity for all stakeholders in the scholarly ecosystem.

In reflecting on this incident, critical questions arise: How might platforms better design their policies and infrastructure to accommodate nuanced realities, such as AI-augmented (but human-directed) research? How can scholars advocate for and practice openness in a way that doesn't inadvertently trigger punitive measures? And how do we preserve the **spirit of open inquiry** while navigating legitimate concerns about new technologies? These questions do not have easy answers, but grappling with them is necessary for academia's progression. The concept of **OntoMotoOS** reminds us that co-creative solutions – ones that integrate multiple perspectives and agents – are possible. Perhaps the resolution in cases like this is not a zero-sum enforcement, but a "win-win" where platforms develop refined rules and scholars adhere to responsible transparency, each side adjusting in light of the other [46]. Our discussion suggested some avenues: policy reform for AI disclosures, improved human-in-the-loop moderation, community-driven standards for AI use in research, and an emphasis on dialogue over summary judgment.

Ultimately, this episode underscores that we are in a **transitional era**. Academic institutions and values are being tested by the advent of AI in much the same way they were tested by the rise of the internet, open access, or data sharing mandates in earlier times. The Figshare case, as painful as it was for the author involved, contributes a cautionary tale that can inform how we move forward. It shows the costs of misalignment between personal scholarly ethics and platform rules, but it also highlights the resilience of the scholarly record (through DOI traces and alternate repositories) and the community's ability to reflect and adapt.

In concluding, it is fitting to strike a note of **appreciation and forward-looking optimism**. While the author experienced disappointment and the loss of a dissemination venue, the incident spurred a deeper inquiry into the very themes that OntoMotoOS is concerned with – governance, transparency, and cooperative problem-solving. In a statement of constructive grace, the author has even expressed gratitude: "This episode may have closed one platform, but it opened another stage: the chance to write, to reflect, and to leave behind a narrative larger than any single DOI. For that paradoxical gift, I can only say — thank you, Figshare." [47] [48]. Such a perspective transforms a setback into an opportunity for growth. It reminds us that in the pursuit of a more ethical and enlightened integration of AI into academia, even conflicts and "scars" [49] can serve as guidance. The hope is that by learning from cases like this, future collaborations between humans, AIs, and institutions will not result in erasures, but in new understanding – ensuring that neither scholarly principles nor institutional integrity have to be casualties of progress.

References

1. Figshare Content Policy (2023). Figshare Terms and Conditions, clause on AI-generated content, which prohibits content "that consists primarily of content generated by artificial intelligence or other automated systems" [28].

2. Kim, Y. (2025a). Toward a World Without War: OntoMotoOS and Philosophical Reinsurance. **Figshare** (disabled) / SSRN (available). DOI: 10.254XX/figshare.xxx (now deactivated). (Case study paper that was flagged on Figshare for containing an AI logo, discussing the OntoMotoOS framework for global peace.)

3. Kim, Y. (2025b). **When Principles Collide: Figshare, AI Logos, and a Win-Win Resolution** [1] [2]. Medium essay (September 2025). (First-person account by the author describing the Figshare incident, the clash between platform policy and transparency principle, and proposing a mutual resolution.)

4. Perkins, M., & Roe, J. (2024). Academic publisher guidelines on AI usage: A ChatGPT-supported thematic analysis. **F1000Research, 12**, 1398 [5] [6]. DOI: 10.12688/f1000research.142411.2. (Study of how major publishers are formulating policies on AI in manuscripts; finds consensus that AI can be used with disclosure, and emphasizes transparency and human responsibility.)

5. Committee on Publication Ethics (COPE) (2023). Guidelines on AI in scholarly writing – e.g., COPE Advice: Authors must be transparent about AI tool usage, disclosing in the Methods section how the AI was utilized [27]. (As cited in Perkins & Roe 2024 analysis.)

6. Kim, Y. (2025c). OntoMotoOS as Philosophical Reinsurance: Between Abstraction and Concreteness. **PhilArchive** (Manuscript) [13]. (Explores OntoMotoOS as a "philosophical reinsurance" framework embedding ethical principles into a meta-operating system; provides conceptual grounding for OntoMotoOS and its modules.)

7. Merton, R. K. (1942). The Normative Structure of Science. **Journal of Legal and Political Sociology, 1**(1). (Classic essay outlining CUDOS norms – Communalism, Universalism, Disinterestedness, Originality, Skepticism – stressing openness and honesty in scientific practice.)

8. **COPE Forum Discussion** (2023). Authorship and AI-tools in scholarly publications. (Available via COPE Forum minutes). (Discussion among editors on how to handle AI contributions; consensus that AI cannot be an author and any use must be declared, to avoid deception in the scholarly record [7] [6].)

9. Gillespie, T. (2018). Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media. **Yale University Press**. (While about social media, provides insight into why platforms enforce rules strictly and the challenges of moderation at scale – relevant to understanding Figshare's perspective on rule enforcement.)

10. Future of Research Repositories Conference (2025). Panel on "AI in Academic Repositories: Threat or Opportunity?". (Hypothetical contemporary source discussing cases like Figshare's and recommending best practices for integrating AI disclosures into repository workflows.)

1 2 9 10 11 12 29 30 31 32 33 34 35 36 37 38 39 40 41 42 46 47 48 49 When Principles Collide: Figshare, AI Logos, and a Win-Win Resolution | by nettalk83 | Sep, 2025 | Medium
https://medium.com/@nettalk83/when-principles-collide-figshare-ai-logos-and-a-win-win-resolution-f03a0a2acdd4

3 4 5 6 7 8 26 27 43 44 45 Academic publisher guidelines on AI usage: A ChatGPT supported thematic analysis - PMC
https://pmc.ncbi.nlm.nih.gov/articles/PMC10844801/

13 17 Toward a World Without War_ OntoMotoOS and Philosophical Reinsurance.pdf
file://file-JGNvZbCkEm1mwwEYke42Rx

14 15 16 18 19 20 21 22 23 24 25 Toward a World Without War_ OntoMotoOS and Philosophical Reinsurance.pdf
file://file-1kkfxxqnPqERSZKSpAZ3qd

28 Figshare terms and conditions
https://figshare.com/terms